Claims

[c1] 1.A method for facilitating writing and retrieval in a content addressed storage (CAS) system utilizing a Logical Unit Number/Logical Block Address (LUN/LBA) interface comprising steps of:

a. receiving an input/output (I/O) operation request, associated application-level LUN/LBA combination, and optionally, content data from an application program at a first LUN/LBA processing tier or at a second OID processing tier,

b. hashing said content data to generate an OID at said second OID processing tier if content data is received in said receiving step,

c. consulting a high-level OID table and an access property with said received application-level LUN/LBA combination or said generated OID to determine whether said content data can be retrieved from or written to an LUN/LBA combination, respectively, and

based on said consulting step, performing at a third storage subsystem LUN/LBA processing tier actions comprising: a write, over-write, or addition operation to an LUN/LBA combination, a retrieval from an LUN/LBA combination, and a denial of operation.

[c2] 2.A method for facilitating writing and retrieval, as per claim 1, wherein said high-level OID table and said low-level OID table are updated if said content data is written to an LUN/LBA combination.

[c3] 3.A method for facilitating writing and retrieval, as per claim 1, wherein said consulting step occurs at said first LUN/LBA processing tier, if said receiving step occurs at said first LUN/LBA processing tier, and at said second OID processing, if said receiving step occurs at said second OID processing tier.

[c4] 4.A method for facilitating writing and retrieval, as per claim 1, wherein said high-level table is comprised of at least: an LUN/LBA combination, and an associated OID or plurality of OIDs.

[c5] 5.A method for facilitating writing and retrieval, as per claim 1, wherein said low-level table is comprised of at least: an LUN/LBA combination, an associated OID or plurality of OIDs, and a counter associated with each OID.

[c6] 6.A method for facilitating writing and retrieval, as per claim 1, wherein said access property is one of a: write-once, write-many, or write-many with versioning property.

[c7] 7.A method for facilitating writing and retrieval, as per claim 6, wherein said write-once property allows a write operation, if an LUN/LBA combination contained in a high-level table has an associated OID with a null value, and a denial of operation otherwise.

[c8] 8.A method for facilitating writing and retrieval, as per claim 6, wherein said write-many property allows operations comprising: a write operation, if an LUN/LBA combination contained in a high-level table has an associated OID with a null value, and a re-write operation, if an LUN/LBA combination contained in a high-level table has an associated OID with a non-null value.

[c9] 9.A method for facilitating writing and retrieval, as per claim 6, wherein said write-many with versioning property allows operations comprising: a write operation, if an LUN/LBA combination contained in a high-level table has an associated OID with a null value, and an addition operation, if an LUN/LBA combination contained in a high-level table has an associated OID or plurality of OIDs with non-null values.

[c10] 10.A method for facilitating writing and retrieval, as per claim 1, wherein said write operation comprises steps of writing an OID generated from a hash of said content

data to either one of, or both, a high-level OID table and a low-level OID table and writing said content data to an LUN/LBA combination.

[c11] 11.A method for facilitating writing and retrieval, as per claim 1, wherein said over-write operation comprises steps of updating an existing OID with an OID generated from a hash of said content data associated with an LUN/LBA combination in either one of or both a high-level OID table and a low-level OID table and writing said content data to an LUN/LBA combination.

[c12] 12.A method for facilitating writing and retrieval, as per claim 1, wherein said addition operation comprises steps of adding to an OID list in either one of or both a high-level OID table and a low-level OID table, an OID generated from a hash of said content data associated with an LUN/LBA combination and writing to an LUN/LBA combination said content data.

[c13] 13.A first LUN/LBA processing tier performing actions comprising:

a. receiving an input/output (I/O) operation request, associated application-level LUN/LBA combination, and optionally, content data from an application program,

b. passing said content data to a second OID processing tier if content data is received from said application pro-

gram,

c. using said application-level LUN/LBA combination received from said application program or using an OID returned from second OID processing tier if said content data was passed to said second OID processing tier in previous step to determine whether corresponding or same OID exists in a high-level OID table,

d. determining whether a write, over-write, or an addition operation is allowed or if an operation is denied, based on said determining step and said access property,

e.passing said operation information and either: an OID if a retrieval operation is requested or an LUN/LBA address combination element location and associated content data for any other operation, and

f. updating said high-level OID table with a new OID for an associated application-level LUN/LBA if a write, over-write, or an addition operation is allowed.

[c14] 14.A second Object ID (OID) processing tier performing actions comprising:

a. receiving an application-level LUN/LBA address combination and optionally, content data from either an application program or a first LUN/LBA processing tier,

b. receiving an OID from a first LUN/LBA processing tier if content data is not received,

c. hashing said content data to generate an OID if content data is received,

d. determining a physical LUN/LBA by using received OID to consult low-level OID table,

e. determining whether an OID exists in a high-level OID table by comparing generated OID to OIDs stored in said high-level OID table,

f. updating high-level OID table with generated OID if generated OID does not exist in a high-level OID table,

g. passing a write request, received content data, and a location of OID in low-level OID table to a third storage subsystem LUN/LBA processing tier if content data is received in receiving step, based on said determining step and operations allowed by said access property,

h. updating a low-level OID table with said generated OID and incrementing a counter associated with an OID, if a write request and associated data is passed to said third storage subsystem LUN/LBA processing tier,

i. passing a retrieval request and a physical LUN/LBA combination to said third storage subsystem LUN/LBA processing tier if an OID was received in receiving step, based on said determining step and operations allowed by said access property,

j. passing to a first LUN/LBA processing tier said OID if a write request and associated data is passed to said third storage subsystem LUN/LBA processing tier, and

k. receiving content data from said third storage subsystem LUN/LBA processing tier, if a retrieval request and said physical LUN/LBA is passed to said third storage subsystem LUN/LBA processing tier.

[c15] 15.An third storage subsystem LUN/LBA processing tier performing actions comprising:
a. receiving from a second OID processing tier a write request or a retrieval request and associated data,
b. using an OID received from said second OID processing tier to determine associated LUN/LBA combination,
c. writing to an LUN/LBA combination if a write request and associated content data is received,
d. retrieving from an LUN/LBA if a retrieval request and an associated OID is received,
e. re-hashing content data retrieved from an LUN/LBA combination if a retrieval request and an associated OID is received,
f. comparing results of re-hashing step with an OID associated with said LUN/LBA combination stored in a low-level OID table, and
g. passing to an second OID processing tier content data if a retrieval request and an associated OID was received.

[c16] 16.A system for CAS having an LUN/LBA interface comprising:
a. a first LUN/LBA processing tier,

b. a second OID processing tier,

c. a storage subsystem LUN/LBA processing,

d. a first table providing a correspondence between a high-level LUN/LBA combination and an OID or plurality of OIDs, and

e. a second table providing a correspondence between a low-level LUN/LBA combination and an OID or plurality of OID.

[c17] 17.A system for CAS having an LUN/LBA interface, as per claim 16, wherein said high-level table is comprised of at least: an LUN/LBA combination, and an associated OID or plurality of OIDs.

[c18] 18.A system for CAS having an LUN/LBA interface, as per claim 16, wherein said low-level table is comprised of at least: an LUN/LBA combination, an associated OID or plurality of OIDs, and a counter associated with each OID.